arXiv:1007.5444v2 [math.NT] 31 Jan 2011

ON CERTAIN OTHER SETS OF INTEGERS

TOM SANDERS

Abstract. We show that if $A \subset \{1, \ldots, N\}$ contains no non-trivial three-term arithmetic progressions then $|A| = O(N/\log^{3/4-o(1)} N)$.

1. Introduction

In this paper we refine the quantitative aspects of Roth's theorem on sets avoiding three-term arithmetic progressions. Our main result is the following.

Theorem 1.1. *Suppose that $A \subset \{1, \ldots, N\}$ contains no non-trivial three-term arithmetic progressions. Then*

$$|A| = O(N/\log^{3/4-o(1)} N).$$

The first non-trivial upper bound on the size of such sets was given by Roth [Rot52, Rot53], and there then followed refinements by Heath-Brown [HB87] and Szemerédi [Sze90], and later Bourgain [Bou99, Bou08], culminating in the above with 2/3 in place of 3/4. Our interest is in bounds so we have not recorded references to the many proofs and generalisations of Roth's theorem not concerned with this issue.

Our argument follows [Bou08] where the key new ingredient was a random sampling technique. We couple this with the heuristic (see [Bou85] for a rigorous formulation) that random sampling 'increases dissociativity' to get a little bit more out of the method. No progress is made on the algebraic analogue in $(\mathbb{Z}/3\mathbb{Z})^n$.

Salem and Spencer [SS42] showed that the surface of high-dimensional convex bodies could be embedded in the integers to construct sets of size $N^{1-o(1)}$ containing no three-term progressions, and Behrend [Beh46] noticed that spheres are a particularly good choice. Recently Elkin [Elk08] tweaked this further by thickening the spheres to produce the largest known progression-free sets, and his argument was then considerably simplified by Green and Wolf in the very short and readable note [GW10b].

In the next section we establish our basic notation before we present a sketch of the main argument in §3. The remainder of the paper then makes the sketch precise.

2. Basic notation: the Fourier transform and Bohr sets

Suppose that G is a finite abelian group. We write $M(G)$ for the space of measures on G and given two measures $\mu, \nu \in M(G)$ define their convolution $\mu * \nu$ to be the measure induced by

$$C(G) \to C(G); f \mapsto \int f(x+y) d\mu(x) d\nu(y).$$

There is a privileged measure $\mu_G \in M(G)$ on G called Haar probability measure, and defined to be the measure assigning mass $|G|^{-1}$ to each element of G.

The significance of this measure is that it is the unique (up to scaling) translation invariant measure on G: for $x \in G$ and $\mu \in M(G)$ we define $\tau_x(\mu)$ to be the measure induced by

$$C(G) \to C(G); f \mapsto \int f(y)d\mu(y+x),$$

and it is easy to see that $\tau_x(\mu_G) = \mu_G$ for all $x \in G$.

We use Haar measure to pass between the notion of function $f \in L^1(\mu_G)$ and measure $\mu \in M(G)$. Indeed, since G is finite we shall often identify μ with $d\mu/d\mu_G$, the Radon-Nikodym derivate of μ with respect to μ_G. In light of this we can extend the notion of translation and convolution to $L^1(\mu_G)$ with ease: given $f \in L^1(\mu_G)$ and $x \in G$ we define the translation of f by x point-wise by

$$\tau_x(f)(y) := \frac{d(\tau_x(fd\mu_G))}{d\mu_G)}(y) = f(x+y) \text{ for all } y \in G,$$

and given $f, g \in L^1(\mu_G)$ we define convolution of f and g point-wise by

$$f * g(x) := \frac{d((fd\mu_G) * (gd\mu_G))}{d\mu_G}(x) = \int f(y)g(x-y)d\mu_G(y),$$

and similarly for the convolution of $f \in L^1(\mu_G)$ with $\mu \in M(G)$.

Convolution operators can be written in a particularly simple form with respect to the Fourier basis which we now recall. We write $\widehat{G}$ for the dual group, that is the finite Abelian group of homomorphisms $\gamma : G \to S^1$, where $S^1 := \{z \in \mathbb{C} : |z| = 1\}$. Given $\mu \in M(G)$ we define $\widehat{\mu} \in \ell^\infty(\widehat{G})$ by

$$\widehat{\mu}(\gamma) := \int \overline{\gamma} d\mu \text{ for all } \gamma \in \widehat{G},$$

and extend this to $f \in L^1(G)$ by $\widehat{f} := \widehat{fd\mu_G}$. It is easy to check that $\widehat{\mu * \nu} = \widehat{\mu} \cdot \widehat{\nu}$ for all $\mu, \nu \in M(G)$ and $\widehat{f * g} = \widehat{f} \cdot \widehat{g}$ for all $f, g \in L^1(\mu_G)$.

It is also useful to have an inverse for the Fourier transform: given $g \in \ell^1(\widehat{G})$ we define $g^\vee \in C(G)$ by

$$g^\vee(x) = \sum_{\gamma \in \widehat{G}} g(\gamma)\gamma(x) \text{ for all } x \in G,$$

and note that $(g^\vee)^\wedge = g$. It is now natural to define the convolution of two functions $f, g \in \ell^1(\widehat{G})$ to be

$$f * g(\gamma) = \sum_{\gamma' \in \widehat{G}} f(\gamma')g(\gamma - \gamma') \text{ for all } \gamma \in \widehat{G},$$

so that $(f * g)^\vee = f^\vee \cdot g^\vee$.

As is typical we shall establish the presence of non-trivial arithmetic progressions by counting them, and to this end we introduce the tri-linear form

$$T(f,g,h) := \int f(x-y)g(y)h(x+y)d\mu_G(x)d\mu_G(y),$$

which has a rather useful Fourier formulation:

$$T(f,g,h) = \sum_{\gamma \in \widehat{G}} \widehat{f}(\gamma)\widehat{g}(-2\gamma)\widehat{h}(\gamma).$$

Following [Bou08] we use a slight generalization of the traditional notion of Bohr set, letting the width parameter vary according to the character. The advantage of this definition is that the meet of two Bohr sets in the lattice of Bohr sets is then just their intersection.

A set B is called a *Bohr set* if there is a *frequency set* Γ of characters on G, and a *width function* $\delta \in (0,2]^\Gamma$ such that

$$B = \{x \in G : |1-\gamma(x)| \leqslant \delta_\gamma \text{ for all } \gamma \in \Gamma\}.$$

The size of the set Γ is called the *rank* of B and is denoted $\mathrm{rk}(B)$. Technically the same Bohr set can be defined by different frequency sets and width functions; we make the standard abuse that when we introduce a Bohr set we are implicitly fixing a frequency set and width function.

There is a natural way of dilating Bohr sets which will be of particular use to us. Given such a B, and $\rho \in \mathbb{R}^+$ we shall write B_ρ for the Bohr set with frequency set Γ and width function[1] $\rho\delta$ so that, in particular, $B = B_1$ and more generally $(B_\rho)_{\rho'} = B_{\rho\rho'}$.

With these dilates we say that a Bohr set B' is a *sub-Bohr set* of another Bohr set B, and write $B' \leqslant B$, if

$$B'_\rho \subset B_\rho \text{ for all } \rho \in \mathbb{R}^+.$$

Finally, we write β_ρ for the probability measure induced on B_ρ by μ_G, and β for β_1.

Throughout the paper Cs will denote absolute, effective, but unspecified constants of size greater than 1 and cs will denote the same of size at most 1. Typically the constants will be subscripted according to the result which they come from and superscripted within arguments.

We shall also find a variant of the big-O notation useful. Specifically, if $X, Y \geqslant 1$ then we write $X = \widetilde{O}(Y)$ to mean $X = O(Y \log^{O(1)} 2Y)$, and if $X, Y \leqslant 1$ then $X = \widetilde{\Omega}(Y)$ to mean $X^{-1} = \widetilde{O}(Y^{-1})$.

3. Sketching the argument

The following discussion is based on the approaches to Roth's theorem developed by Bourgain in [Bou99] and then [Bou08]. There have been a number of expositions of the first of these papers – the interested reader may wish to consult Tao and Vu [TV06] – and the second paper contains quite a detailed heuristic discussion of the argument.

---

[1] Technically width function $\gamma \mapsto \min\{\rho\delta_\gamma, 2\}$.

It should be remarked that the present section is not logically necessary for the rest of the paper, and the reader who is not already familiar with the basic ideas may find it difficult to follow.

We shall concentrate on the problem of showing that if G is a finite abelian group of odd order and $A \subset G$ has density $\alpha \in (0,1]$ then

$$T(1_A, 1_A, 1_A) = \exp(-\widetilde{O}(\alpha^{-C})),$$

ideally with as small a constant C as possible.

It seems unlikely that the odd order condition is necessary in light of the work of Lev [Lev04], but removing it seems to add technicalities which obscure the underlying ideas. In any case the passage from finite abelian groups of odd order to $\{1, \dots, N\}$ is unaffected and follows from the usual Freĭman embedding.

To begin recall that by the usual application of the inversion formula and the triangle inequality, if A has somewhat fewer than expected three-term progressions then

$$\sum_{0_{\widehat{G}} \neq \gamma \in \mathrm{Spec}_{\Omega(\alpha)}(1_A)} |\widehat{1_A}(\gamma)|^2 |\widehat{1_A}(-2\gamma)| = \Omega(\alpha^3). \tag{3.1}$$

3.1. **The basic approach.**

Roth [Rot53] noted that the above implies, in particular, that the large spectrum is non-empty, and so we have a non-trivial character γ with $|\widehat{1_A}(\gamma)| = \Omega(\alpha^2)$. This can be used to give a density increment on (a translate of) $I := \{x \in G : \gamma(x) \approx 1\}$ of the form $\alpha \mapsto \alpha(1 + \Omega(\alpha))$.

In [Bou99] Bourgain noted that approximate level sets of characters – sets like I – behave roughly like groups, and that crucially the preceding argument can be run relative to anything which behaves in that way.

The intersection of many approximate level sets is a Bohr set (hence their importance to us) and working through the previous for a set A of relative density α on a d-dimensional Bohr set B gives a new Bohr set B' with

$$d \mapsto d + 1 \text{ and } \alpha \mapsto \alpha(1 + \Omega(\alpha)). \tag{3.2}$$

In general we only have a bound of the form

$$\mu_G(B') \geqslant \exp(-\widetilde{O}(d))\mu_G(B), \tag{3.3}$$

although for groups of bounded exponent this can be much improved. Indeed, if G is the model group $(\mathbb{Z}/3\mathbb{Z})^n$ then $\mu_G(B') = \Omega(\mu_G(B))$, and this is the reason for the difference in the bounds in the model setting (see [Mes95]) and here.

A density increment of the form in (3.2) cannot proceed for more that $\widetilde{O}(\alpha^{-1})$ steps, and so there must be some point where A has about the expected number of three-term progressions on a Bohr set B' of dimension $d + \widetilde{O}(\alpha^{-1})$. It follows that

$$\mu_G(B') \geqslant \exp(-\widetilde{O}(d)) \dots \exp(-\widetilde{O}(d + \alpha^{-1}))\mu_G(B)$$

which tells us that

$$T(1_A, 1_A, 1_A) = \exp(-\widetilde{O}(d\alpha^{-1} + \alpha^{-2}))\mu_G(B)^2. \tag{3.4}$$

We are interested in the case $d = 0$ and $B = G$ when this gives the bound in [Bou99].

3.2. **The energy increment approach.** The previous argument only uses the fact that (3.1) tells us the large spectrum is non-empty. It is natural to try to get a density increment by noting that the Bohr set B with frequency set $\mathrm{Spec}_{\Omega(\alpha)}(1_A)$ majorises that spectrum, and hence $1_A * \beta$ has large $L^2(\mu_G)$-mass.

An energy increment technique motivated by the work of Heath-Brown [HB87] and Szemerédi [Sze90] (essentially Lemma 7.2) now tells us that A has density $\alpha(1 + \Omega(1))$ on (a translate of) B. It remains to estimate the dimension of B. This is certainly at most the size of the largest dissociated subset of the large spectrum, and this can be bounded by Chang's theorem [Cha02]; in this case it is $\widetilde{O}(\alpha^{-2})$.

All of this can, itself, be done for a set A of relative density α in a Bohr set B of dimension d. In which case either we have many three-term progressions or else we have a new Bohr set with

$$d \mapsto d + \widetilde{O}(\alpha^{-2}) \text{ and } \alpha \mapsto \alpha(1 + \Omega(1)).$$

This process can be iterated at most $O(\log 2\alpha^{-1})$ times and so proceeding as before with the lower bound (3.3) we conclude that

$$T(1_A, 1_A, 1_A) = \exp(-\widetilde{O}(d + \alpha^{-2}))\mu_G(B)^2. \tag{3.5}$$

This represents no improvement over (3.4) in our case of interest, but for large values of d it does and we shall find this fact useful later on.

3.3. **Comparing the approaches.** In the approach in §3.1 we only looked at $\widetilde{O}(\alpha^{-1})$ different characters but we did so one at a time and every time we iterated we used (3.3) which cost us a factor exponential in the dimension of the Bohr set. In the approach in §3.2 we looked at $\widetilde{O}(\alpha^{-2})$ characters but took almost all of them together meaning that we only lost the factor from (3.3) a very small number of times.

To get a better bound one should like to improve Chang's theorem, but that is not possible without additional information about A. (See [Gre03] for details.) However, suppose that we are in the extreme case of Chang's theorem and there are many dissociated characters in the large spectrum. They each lead to a different density increment of the type in §3.1 and one might hope that these increments are in some sense 'independent' (since the characters are) so that rather than having to do them one after the other and suffering the expense of repeated application of (3.3), we could do them all at once and only suffer the cost once. It is this hope to which we now turn our attention following [Bou08].

3.4. **The random sampling strategy.** In the first instance the argument proceeds by removing the mass in (3.1) corresponding to $\mathrm{Spec}_{\sqrt{\alpha}}(A)$ using the Heath-Brown-Szemerédi energy increment technique as in §3.2. In this case we get a Bohr set of dimension $\widetilde{O}(\alpha^{-1})$ on which we have a density increment of the form $\alpha \mapsto \alpha(1 + \Omega(1))$. Otherwise, by dyadic decomposition, there is some $\tau \in [\Omega(\alpha), \sqrt{\alpha}]$ such that

$$\sum_{\gamma \in S_\tau} |\widehat{1_A}(\gamma)|^2 = \widetilde{\Omega}(\tau^{-1}\alpha^2), \tag{3.6}$$

where $S_\tau := \mathrm{Spec}_{2\tau}(1_A) \setminus \mathrm{Spec}_\tau(1_A)$. We shall split into two cases depending on whether S_τ contains a large dissociated set or not.

Suppose that T is a dissociated subset of S_τ, and let c_λ be the phase of $\widehat{1_A}(\lambda)$. By expanding out we have

$$\langle 1_A, \prod_{\lambda \in T} (1 + c_\lambda \operatorname{Re} \lambda)\rangle_{L^2(\mu_G)} = \alpha + \sum_{\lambda \in \Lambda} |\widehat{1_A}(\lambda)| + \dots . \tag{3.7}$$

If the remaining terms really were an error term then we would have that the inner product is at least $\alpha(1 + \Omega(\tau|T|))$.

Now, a Riesz product p generated by a set Λ of characters is essentially invariant under convolution with β, where B is the Bohr set with frequency set Λ. Thus

$$\langle 1_A, p\rangle_{L^2(\mu_G)} \approx \langle 1_A * \beta, p\rangle_{L^2(\mu_G)} \leqslant \|1_A * \beta\|_{L^\infty(\mu_G)} \int p d\mu_G.$$

Of course, $1_A * \beta(x)$ is just the density of A on $x + B$ and so if the inner product on the left is large and the integral of the Riesz product is 1 then we have a density increment.

Since T is dissociated, the integral of our Riesz product is 1 and we would get the density increment $\alpha \mapsto \alpha(1 + \Omega(\tau|T|))$ on a Bohr set of dimension $O(|T|)$ if the remaining terms in (3.7) were error terms. Of course, they are not and to deal with this we sample from T. Let Λ be a subset of T with $\lambda \in T$ chosen independently with probability θ. Then

$$\begin{aligned} \mathbb{E}\langle 1_A, \prod_{\lambda \in \Lambda} (1 + c_\lambda \operatorname{Re} \lambda)\rangle_{L^2(\mu_G)} &= \langle 1_A, \prod_{\lambda \in T} (1 + \theta c_\lambda \operatorname{Re} \lambda)\rangle_{L^2(\mu_G)} \\ &\geqslant \langle 1_A, \exp(\sum_{\lambda \in T} \theta c_\lambda \operatorname{Re} \lambda)\rangle_{L^2(\mu_G)}(1 + O(\theta^2|T|)) \\ &\geqslant \langle 1_A, 1 + \sum_{\lambda \in T} \theta c_\lambda \operatorname{Re} \lambda\rangle_{L^2(\mu_G)}(1 + O(\theta^2|T|)) \\ &\geqslant \alpha + \theta|T|\tau\alpha/2 \end{aligned} \tag{3.8}$$

if $\theta \leqslant c\tau$. Now we get a Bohr set of (expected) dimension $O(\theta|T|)$ on which we have density $\alpha(1 + \Omega(\tau\theta|T|))$.

This can all be done for a set A of relative density α in a d-dimensional Bohr set, and taking $\theta = c\tau$ we have a dichotomy: if there are at least k dissociated elements in S_τ then we get a new Bohr set with

$$d \mapsto d + O(\tau k) \text{ and } \alpha \mapsto \alpha(1 + \Omega(\tau^2 k)); \tag{3.9}$$

if there are at most k dissociated elements in S_τ then we apply the Heath-Brown-Szemerédi energy increment technique to get a Bohr set of dimension $d + O(k)$ on which we have density $\widetilde{\Omega}(\tau^{-1}\alpha)$ by equation (3.6). In this instance we apply (3.5) and are done.

Now, (3.9) can happen at most $O(\tau^{-2}k^{-1})$ times which means that we have at most that many applications of (3.3), while the dimension of the Bohr set is at most

$$O(\tau^{-2}k^{-1}).O(\tau k)+\widetilde{O}(\alpha^{-1})=\widetilde{O}(\alpha^{-1})$$

since $\tau=\Omega(\alpha)$. After that the iteration must terminate either with many progressions or an application of (3.5) and we find that for $A\subset G$ of density α we have

$$T(1_A,1_A,1_A)\geqslant\exp(-\widetilde{O}(\tau^{-2}k^{-1}\alpha^{-1}+k+\tau^2\alpha^{-2})). \tag{3.10}$$

Of course $\tau\in[\Omega(\alpha),\sqrt{\alpha}]$ so we can optimise with $k=\alpha^{-3/2}$ to get the main result in [Bou08].

3.5. **Relaxing the dissociation.** To get the density increment from (3.8) we used the fact that T is dissociated so that the Riesz product has integral 1, but if the integral is at most $1+c\theta|T|\tau$ then the argument we used still works. Moreover, the random sampling process actually makes the integral closer to 1, and it is this which we shall exploit to improve (3.10).

We start with a decomposition of Bourgain from [Bou90] the existence of which is an easy induction: for a parameter $m\in\mathbb{N}$,

(i) either half of S_τ has no dissociated set of size at least m;
(ii) or half of S_τ is a disjoint union of dissociated sets of size m.

In the former case we get a Bohr set of dimension m on which we have density $\Omega(\tau^{-1})=\Omega(\alpha^{1/2})$ by the Heath-Brown-Szemerédi energy increment technique, and we are done after applying the bound (3.5). In the latter case write T for the union, and note by Rudin's inequality and the triangle inequality that

$$1_T*\cdots*1_T(0_{\widehat{G}})\leqslant O(r/m)^{(r-2)/2}|T|^{r-1}, \tag{3.11}$$

where the convolution is r-fold. This sort of expression and related consequences have been exploited extensively in the work of Shkredov [Shk08b, Shk09] and Shkredov and Yekhanin [SY10], and that work was one of the starting points for this paper.

Let Λ be a subset of T with $\lambda\in T$ chosen independently with probability θ. Now, suppose (for simplicity) that all the phases c_λ from the previous subsection are 1. Then we have

$$\begin{aligned}\mathbb{E}\int\prod_{\lambda\in\Lambda}(1+c_\lambda\operatorname{Re}\lambda)d\mu_G &\leqslant 1+\theta\int\sum_{\lambda\in\Lambda}c_\lambda\operatorname{Re}\lambda d\mu_G+\sum_{r=2}^{\infty}\frac{\theta^r}{r!}1_T*\cdots*1_T(0_{\widehat{G}})\\ &\leqslant 1+\sum_{r=2}^{\infty}\frac{m}{r|T|}O\left(\frac{\theta|T|}{\sqrt{mr}}\right)^r\end{aligned}$$

since none of the characters is trivial. Of course if $|T|\geqslant C\tau^{-1}\sqrt{m}$ then we may take $\theta|T|=c'\sqrt{m}$ whilst keeping $\theta\leqslant c\tau$, and the upper bound is then just $1+O(\theta^2|T|)$. Assuming that this typical behaviour actually happens we get a density increment of $\alpha\mapsto\alpha(1+\Omega(\tau\sqrt{m}))$ on a Bohr set of dimension $O(\sqrt{m})$. For a suitable choice of m this will be an improvement on what we had before.

In the non-relative situation it is easy to see that T is as large as needed since it's at least half of S_τ, which in turn has size $\Omega(\tau^{-3})$ from (3.6). However, when we relativise to Bohr sets this is not true any more. To solve this, if $|T| \leqslant C\tau^{-1}\sqrt{m}$ then we shall arrange things so that we can take a subset T' of size $\eta|T|$ such that

$$\sum_{\lambda \in T'} |\widehat{1_A}(\gamma)|^2 = \widetilde{\Omega}(\eta\tau^{-1}\alpha^2),$$

and apply the usual Heath-Brown-Szemerédi energy increment technique to get a Bohr set of dimension $O(\eta\tau^{-1}\sqrt{m})$ on which we have density $\widetilde{\Omega}(\eta\tau^{-1}\alpha)$; we finish off this case when $|T|$ is small by applying (3.5).

Relativising all this as in the previous subsection we get that our iteration procedure terminates inside $O(\tau^{-1}m^{-1/2})$ steps and

$$T(1_A, 1_A, 1_A) \geqslant \exp(-\widetilde{O}(\tau^{-1}m^{-1/2}\alpha^{-1} + m + \eta\tau^{-1}\sqrt{m} + \eta^{-2}\tau^2\alpha^{-2})).$$

Of course $\tau \in [\Omega(\alpha), \sqrt{\alpha}]$ so we can optimise with[2] $\eta = \tau\alpha^{-1/2}$ and $m = \alpha^{-4/3}$ to get our main theorem.

4. The large spectrum

In this section we shall be working with respect to a probability measure μ on G. Later we shall take μ to be a sort of approximate Haar measure on approximate groups, but for the moment our analysis does not require this.

Our object of study is the spectrum of functions with respect to μ: given a function $f \in L^1(\mu)$ and a parameter $\epsilon \in (0, 1]$ we define the *ϵ-spectrum of f w.r.t. μ* to be the set

$$\mathrm{Spec}_\epsilon(f, \mu) := \{\gamma \in \widehat{G} : |(f d\mu)^\wedge(\gamma)| \geqslant \epsilon\|f\|_{L^1(\mu)}\}.$$

This definition extends the usual one from the case $\mu = \mu_G$. In that case there is a rather simple estimate for the size of the spectrum which follows from Bessel's inequality, and it will be useful for us to have a relative version of this.

Defining notions by what they do, rather than what they are is a fairly standard idea, but it has seen particular success in the development of approximate analogues of concepts in additive combinatorics. This is spelt out particularly clearly by Gowers and Wolf in [GW10a], and the next definition (as well as similar later ones) is motivated by their discussion of quadratic rank. We say that $\Lambda \subset \widehat{G}$ is *K-orthogonal w.r.t. μ* if

$$\int |1 + g^\vee|^2 d\mu \leqslant (1 + K)(1 + \|g\|^2_{\ell^2(\Lambda)}) \text{ for all } g \in \ell^2(\Lambda).$$

Lemma 4.1 (Monotonicity of orthogonality)**.** *Suppose that μ' is another probability measure, Λ is K-orthogonal w.r.t. μ, $\Lambda' \subset \Lambda$ and $K' \geqslant K$. Then Λ' is K'-orthogonal w.r.t. $\mu' * \mu$.*

---

[2] There is some flexibility with η; m is the critical parameter.

Proof. The only part worthy of remark is the replacing of μ by $\mu' * \mu$: this is immediate on noting that $\tau_{-y}(1+g^\vee) = 1 + h^\vee$ where $h(\lambda) = \overline{\lambda(y)}g(\lambda)$ and $\|h\|_{\ell^2(\Lambda)}^2 = \|g\|_{\ell^2(\Lambda)}^2$. It follows that Λ is K-orthogonal w.r.t. $\tau_y(\mu)$, and we get the result on integrating. □

We define the (K,μ)*-relative size* of Λ to be the size of the largest subset of Λ that is (K,μ)-orthogonal w.r.t. μ. By monotonicity the $(K', \mu' * \mu)$-relative size dominates the (K,μ)-relative size; the $(0,\mu_G)$-relative size of Λ is the size of the largest subset of Λ not containing $\{0_{\widehat{G}}\}$.

The Bessel bound mentioned earlier follows easily from this definition.

Lemma 4.2 (The Bessel bound)**.** *Suppose that* $f \in L^2(\mu)$ *is not identically zero and write* $L_f := \|f\|_{L^2(\mu)}\|f\|_{L^1(\mu)}^{-1}$. *Then* $\mathrm{Spec}_\epsilon(f,\mu)$ *has* $(1,\mu)$*-relative size* $O(\epsilon^{-2}L_f^{-2})$.

Proof. Suppose that $\Lambda \subset \mathrm{Spec}_\epsilon(f,\mu)$ is $(1,\mu)$-orthogonal. Then

$$1 + \|g^\vee\|_{L^2(\mu)}^2 = \frac{1}{2}\left(\int |1-g^\vee|^2 d\mu + \int |1+g^\vee|^2 d\mu\right) \leqslant 2(1 + \|g\|_{\ell^2(\Lambda)}^2)$$

for all $g \in \ell^2(\Lambda)$. It follows (by letting $\|g\|_{\ell^2(\Lambda)} \to \infty$) that the map $\ell^2(\Lambda) \to L^2(\mu)$ has norm at most $2^{1/2}$ and so by duality the map $L^2(\mu) \to \ell^2(\Lambda)$ defined by $f \mapsto (fd\mu)^\wedge|_\Lambda$ has norm at most $2^{1/2}$, whence

$$|\Lambda|\epsilon^2\|f\|_{L^1(\mu)}^2 \leqslant \sum_{\lambda\in\Lambda} |(fd\mu)^\wedge(\lambda)|^2 \leqslant 2\|f\|_{L^2(\mu)}^2.$$

We get the result on rearranging. □

In the case $\mu = \mu_G$ this yields the usual Parseval bound which pervades work in the area. The bound was first understood in the context of Bessel's inequality in the paper [GT08] of Green and Tao.

Orthogonal random variables can be thought of as roughly pair-wise independent, and naturally we can improve the above bound if we insist on a greater level of independence. The idea of doing this was developed by Chang in [Cha02] following on from work of Bourgain [Bou90].

Given a set of characters Λ and a function $\omega : \Lambda \to D := \{z \in \mathbb{C} : |z| \leqslant 1\}$ we define

$$p_{\omega,\Lambda} := \prod_{\lambda\in\Lambda}(1 + \mathrm{Re}(\omega(\lambda)\lambda)),$$

and call such a function a *Riesz product for* Λ. It is easy to see that all Riesz products are real non-negative functions. They are at their most useful when they also have mass close to 1: the set Λ is said to be K*-dissociated w.r.t.* μ if

$$\int p_{\omega,\Lambda} d\mu \leqslant \exp(K) \text{ for all } \omega : \Lambda \to D.$$

If μ is positive definite then one sees by Plancherel's theorem that the maximum above is attained for the constant function 1 and so we have the following lemma.

Lemma 4.3 (Dissociativity for positive definite measures)**.** *Suppose that* Λ *is a set of characters and* μ *is positive definite. Then* Λ *is* $\log\int p_{1,\Lambda}d\mu$*-dissociated w.r.t.* μ*.*

As before we also have a monotonicity result.

Lemma 4.4 (Monotonicity of dissociativity)**.** *Suppose that* μ' *is another probability measure,* Λ *is* K*-dissociated w.r.t.* μ*,* $\Lambda' \subset \Lambda$ *and* $K' \geqslant K$*. Then* Λ' *is* K'*-dissociated w.r.t.* $\mu' * \mu$*.*

The (K,μ)*-relative entropy* of a set Λ is the size of the largest subset of Λ that is K-dissociated w.r.t. μ. As before relative entropy is monotone; the $(0,\mu_G)$-relative entropy of a set is the size of its largest dissociated (in the usual sense) subset.

A slightly harder property of dissociated sets is that they satisfy a Chernoff-type estimate. The proof of this that we give here essentially localises [GR07, Proposition 3.4].

Lemma 4.5. *Suppose that* Λ *is* K*-dissociated w.r.t.* μ*. Then*

$$\int |\exp(g^{\vee})|d\mu \leqslant \exp(K + \|g\|_{\ell^2(\Lambda)}^2/2) \text{ for all } g \in \ell^2(\Lambda).$$

Proof. Suppose that $g \in \ell^2(\Lambda)$ and begin by noting that

$$\int |\exp(g^{\vee})|d\mu = \int \exp(\operatorname{Re} g^{\vee})d\mu = \int \prod_{\lambda\in\Lambda} \exp(\operatorname{Re}(g(\lambda)\lambda))d\mu.$$

Now we have the elementary inequality $\exp(ty) \leqslant \cosh t + y\sinh t$ whenever $t \in \mathbb{R}$ and $-1 \leqslant y \leqslant 1$, so

$$\int |\exp(g^{\vee})|d\mu \leqslant \int \prod_{\lambda\in\Lambda: g(\lambda)\neq 0} \left(\cosh|g(\lambda)| + \frac{\operatorname{Re}(g(\lambda)\lambda)}{|g(\lambda)|}\sinh|g(\lambda)|\right)d\mu,$$

with the usual convention that $t^{-1}\sinh t$ is 1 if $t = 0$. We define $\omega \in \ell^\infty(\Lambda)$ by

$$\omega(\lambda) := \frac{g(\lambda)\sinh|g(\lambda)|}{|g(\lambda)|\cosh|g(\lambda)|}$$

so that we certainly have $\|\omega\|_{\ell^\infty(\Lambda)} \leqslant 1$, and hence $p_{\omega,\Lambda}$ is a Riesz product and

$$\int |\exp(g^{\vee})|d\mu \leqslant \prod_{\lambda\in\Lambda}\cosh|g(\lambda)| \int p_{\omega,\Lambda}d\mu.$$

The result follows since $\cosh x \leqslant \exp(x^2/2)$. □

It is, perhaps, instructive to compare the conclusion of this lemma with the definition of being K-orthogonal w.r.t. μ.

Green and Ruzsa used the above estimate (for $K = 0$) in [GR07] to prove Chang's theorem for sets, but our argument will proceed along the more traditional lines of establishing Rudin's inequality (w.r.t. to μ) and then applying duality so that it applies to general functions.

Lemma 4.6 (The Chang bound). *Suppose that $f \in L^2(\mu)$ is not identically zero and write $L_f := \|f\|_{L^2(\mu)}\|f\|_{L^1(\mu)}^{-1}$. Then $\mathrm{Spec}_\epsilon(f,\mu)$ has $(1,\mu)$-relative entropy $O(\epsilon^{-2}\log 2L_f)$.*

Proof. Suppose that $\Lambda \subset \mathrm{Spec}_\epsilon(f,\mu)$ is K-dissociated w.r.t. μ. First we establish a relative version of Rudin's inequality. Suppose that $k \in \mathbb{N}$ and $g \in \ell^2(\Lambda)$ has $\|g\|_{\ell^2(\Lambda)} = 2\sqrt{k}$. Then Lemma 4.5 tells us that

$$\|\operatorname{Re} g^\vee\|_{L^{2k}(\mu)} \leqslant (2k)!^{1/2k}\exp(1/2k + 4k/4k) = O(\sqrt{k}\|g\|_{\ell^2(\Lambda)}).$$

However, since $\operatorname{Re}((-i)g^\vee) = \operatorname{Im} g^\vee$ we conclude that $\|g^\vee\|_{L^{2k}(\mu)} = O(\sqrt{k}\|g\|_{\ell^2(\Lambda)})$ and hence that the map $g \mapsto g^\vee$ from $\ell^2(\Lambda)$ to $L^{2k}(\mu)$ has norm $O(\sqrt{k})$.

Now, as with the Bessel bound, by duality and convexity of the $L^p(\mu)$-norms we see that

$$|\Lambda|\epsilon^2\|f\|_{L^1(\mu)}^2 \leqslant \sum_{\lambda\in\Lambda}|(f d\mu)^\wedge(\lambda)|^2 = O(k\|f\|_{L^{2k/(2k-1)}(\mu)}^2) = O(k\|f\|_{L^1(\mu)}^2 L_f^{2/k}).$$

The result follows on putting $k = \lceil \log 2L_f \rceil$ and rearranging. □

Although Chang's theorem cannot be significantly improved, there are some small refinements and discussions of their limitations in the work [Shk06, Shk07] and [Shk08a] of Shkredov. Our approach is largely insensitive to these developments and so we shall not concern ourselves with them here.

Instead of trying to improve Chang's theorem we are going to show that if the large spectrum of a function contains many dissociated elements then we get correlation with a Riesz product. The result should be compared with Proposition (*) from [Bou08] which has a very readable proof in that paper.

On a first reading one may wish to take $\mu = \mu'$. When we use the lemma we shall take μ' positive definite and such that $\mu * \mu' \approx \mu$. It then follows that the integral on the right of the second conclusion is roughly an upper bound for the constant of dissociativity of Λ''' w.r.t. μ. (The reader may wish to look forwards to Lemma 7.3 to get some idea of what we need the lemma below to dovetail with.)

Lemma 4.7. *Suppose μ' is another probability measure, A is an event with $\alpha := \mu(A) > 0$, $\Lambda \subset \mathrm{Spec}_\tau(1_A,\mu)$ is η-orthogonal w.r.t. μ' and has size k, and $m \geqslant 1$ is a parameter. Then at least one of the following is true:*

(i) *there is a set Λ' of $(1,\mu')$-relative entropy $O(m\log 2k)$ which contains at least $k/2$ of the elements of Λ;*

(ii) *for any d with $C_{4.7}\log 2\tau^{-1} \leqslant d \leqslant c_{4.7}\min\{\tau k, \sqrt{m}\}$ there is some set Λ''' of $O(d)$ characters and function $\omega' : \Lambda''' \to D$ such that*

$$\langle 1_A, p_{\omega',\Lambda'''}\rangle_{L^2(\mu)} \geqslant \alpha(1+\tau d/4)\int p_{1,\Lambda'''}d\mu';$$

provided $\eta \leqslant c'_{4.7}\tau$.

Proof. Let $l := \lceil 2m\log 2k\rceil$ and decompose Λ as a disjoint union of sets $\Lambda', \Lambda_1, \ldots, \Lambda_r$ where

(i) any subset of Λ' that is 1-dissociated w.r.t. μ' has size at most l;
(ii) every Λ_i has size l and is 1-dissociated w.r.t. μ'.

If $|\Lambda'|$ is at least $k/2$ then we are in the first case of the lemma; thus we may assume that $\Lambda'' := \Lambda_1 \sqcup \cdots \sqcup \Lambda_r$ has size at least $k/2$.

Let $\omega : \Lambda'' \rightarrow D$ be such that

$$\omega(\lambda)(1_A d\mu)^\wedge(\lambda) = |(1_A d\mu)^\wedge(\lambda)| \text{ for all } \lambda \in \Lambda'';$$

and let $(X_\lambda)_{\lambda\in\Lambda''}$ be independent identically distributed random variables with

$$\mathbb{P}(X_\lambda = 1) = \theta \text{ and } \mathbb{P}(X_\lambda = 0) = 1 - \theta, \text{ where } \theta := d/k,$$

possible by assuming $d \leqslant k$. Since the X_λs are independent we have that

$$\mathbb{E}p_{\omega'' X,\Lambda''}(x) = \prod_{\lambda\in\Lambda''} \mathbb{E}(1 + X_\lambda \operatorname{Re}(\omega''(\lambda)\lambda(x))) = \prod_{\lambda\in\Lambda''} (1 + \theta \operatorname{Re}(\omega''(\lambda)\lambda(x))), \tag{4.1}$$

for any $\omega'' : \Lambda'' \rightarrow D$. We write $\mathcal{N}$ for the event that at most $2ed$ of the variables are 1 and specialise to $\omega'' = \omega$ so that again, by independence,

$$\begin{aligned}
\mathbb{E}1_{\mathcal{N}} p_{\omega X,\Lambda''}(x) &= \mathbb{E}p_{\omega X,\Lambda''}(x) - \sum_{r>2ed} \theta^r \sum_{S\subset\Lambda'':|S|=r} \prod_{\lambda\in S} \operatorname{Re}(\omega(\lambda)\lambda(x))\\
&\geqslant \mathbb{E}p_{\omega X,\Lambda''}(x) - \sum_{r>2ed} \theta^r \binom{|\Lambda''|}{r}\\
&\geqslant \mathbb{E}p_{\omega X,\Lambda''}(x) - \exp(-\Omega(d))
\end{aligned}$$

since $\binom{n}{m} \leqslant (en/m)^m$ for all non-negative integers n and m.

Of course, $1 + z \geqslant \exp(z - 2z^2)$ whenever $-1/2 \leqslant z \leqslant 1/2$, and $\exp(z) \geqslant 1 + z$ for all z so we have (since we may assume $\theta \leqslant 1/2$) that

$$\begin{aligned}
\mathbb{E}p_{\omega X,\Lambda''}(x) &\geqslant \prod_{\lambda\in\Lambda''} \exp(\theta \operatorname{Re}(\omega(\lambda)\lambda(x)) - 2\theta^2)\\
&= \exp(-2\theta^2 k)\exp(\theta \operatorname{Re} \omega^\vee(x)) \geqslant \exp(-2\theta^2 k)(1 + \theta \operatorname{Re} \omega^\vee(x)).
\end{aligned}$$

Coupled with non-negativity of 1_A this means that

$$\begin{aligned}
\mathbb{E}1_{\mathcal{N}}\langle 1_A, p_{\omega X,\Lambda''}\rangle_{L^2(\mu)} &= \langle 1_A, \mathbb{E}p_{\omega X,\Lambda''}\rangle_{L^2(\mu)} - \alpha\exp(-\Omega(d))\\
&\geqslant \exp(-2\theta^2 k)(\alpha + \theta \sum_{\lambda\in\Lambda''} |(1_A d\mu)^\wedge(\lambda)|) - \alpha\exp(-\Omega(d))\\
&\geqslant (1 - 2\theta^2 k)(\alpha + \theta\tau\alpha k/2) - \alpha\exp(-\Omega(d)). \qquad (4.2)
\end{aligned}$$

In the other direction we specialise to $\omega'' \equiv 1$ in (4.1) and get

$$\mathbb{E}p_{X,\Lambda''}(x) = \prod_{\lambda\in\Lambda''} (1 + \theta \operatorname{Re} \lambda(x)) \leqslant \exp(\theta \sum_{\lambda\in\Lambda''} \operatorname{Re} \lambda(x)) = \exp(\theta \operatorname{Re} 1_{\Lambda''}^\vee(x)).$$

However, $\exp(z) \leqslant 1 + z + z^2\exp(z)$ for all $z \in \mathbb{R}$ so that

$$\mathbb{E}\int p_{X,\Lambda''} d\mu' \leqslant 1 + \int \theta \operatorname{Re} 1_{\Lambda''}^\vee d\mu' + \int (\theta \operatorname{Re} 1_{\Lambda''}^\vee)^2 \exp(\theta \operatorname{Re} 1_{\Lambda''}^\vee) d\mu'. \tag{4.3}$$

The second term on the right is estimated using the η-orthogonality of Λ. Indeed, suppose that $|z| = 1$, then

$$\begin{aligned} 2z\int \operatorname{Re} 1_{\Lambda''}^{\vee} d\mu' &= \sum_{\lambda\in\Lambda''} 2\operatorname{Re}\int z\lambda d\mu' \\ &= \sum_{\lambda\in\Lambda''}\left(\int |1+z\lambda|^2 d\mu' - 2\right) \leqslant 2\eta k \end{aligned}$$

since $|\Lambda''| \leqslant k$. Thus, on suitable choice of z, we see that the relevant term in (4.3) is at most ηk in absolute value. To estimate the third term on the right in (4.3) we apply Hölder's inequality with index $q := 1 + \log k$ and conjugate index q' to get that

$$\int (\theta \operatorname{Re} 1_{\Lambda''}^{\vee})^2 \exp(\theta \operatorname{Re} 1_{\Lambda''}^{\vee}) d\mu' \leqslant \theta^2 \left(\int |\operatorname{Re} 1_{\Lambda''}^{\vee}|^{2q'} d\mu'\right)^{1/q'} \left(\int \exp(\theta q \operatorname{Re} 1_{\Lambda''}^{\vee}) d\mu'\right)^{1/q}.$$

Now $|\operatorname{Re} 1_{\Lambda''}^{\vee}| \leqslant k$ and

$$\int |\operatorname{Re} 1_{\Lambda''}^{\vee}|^2 d\mu' \leqslant \int |1_{\Lambda''}^{\vee}|^2 d\mu' \leqslant 2k$$

since Λ is certainly 1-orthogonal. It follows that

$$\left(\int |\operatorname{Re} 1_{\Lambda''}^{\vee}|^{2q'} d\mu'\right)^{1/q'} \leqslant ((2k).k^{2q'-2})^{1/q'} = O(k),$$

and combining what we have so far gives

$$\mathbb{E}\int p_{X,\Lambda''} d\mu' \leqslant 1 + O(\eta\theta k) + O\left(\theta^2 k \left(\int \exp(\theta q \operatorname{Re} 1_{\Lambda''}^{\vee}) d\mu'\right)^{1/q}\right).$$

Recalling that each Λ_i is 1-dissociated w.r.t. μ' and that Λ'' is their disjoint union it follows from Lemma 4.5 that

$$\begin{aligned} \int \exp(\theta q \operatorname{Re} 1_{\Lambda''}^{\vee}) d\mu' &\leqslant \prod_{i=1}^{r}\left(\int \exp(rq\theta \operatorname{Re} 1_{\Lambda_i}^{\vee}) d\mu'\right)^{1/r} \\ &\leqslant \prod_{i=1}^{r} \exp(1/r)\exp(rq^2\theta^2 \|1_{\Lambda_i}\|_{\ell^2(\Lambda_i)}^2/2). \end{aligned}$$

Now, $\|1_{\Lambda_i}\|_{\ell^2(\Lambda_i)}^2 = l$ and $rl \leqslant k$, so

$$\mathbb{E}\int p_{X,\Lambda''} d\mu' \leqslant 1 + O(\theta\eta k) + O(\theta^2 k \exp(k^2 q\theta^2/2l)).$$

Inserting this into (4.2) we get that

$$\begin{aligned} \mathbb{E}1_{\mathcal{N}}\langle 1_A, p_{\omega X,\Lambda''}\rangle_{L^2(\mu)} &\geqslant \alpha\left((1 - O(\eta d) - O(d^2k^{-1}\exp(d^2q/2l)))(1+\tau d/2)\right. \\ &\quad \left. - \exp(-\Omega(d))\right)\mathbb{E}\int p_{X,\Lambda''} d\mu'. \end{aligned}$$

It follows that there are absolute constants c, c' and C such that if

$$\eta \leqslant c\tau,\ dk^{-1}\exp(d^2q/2l) \leqslant c'\tau \text{ and } d \geqslant C\log 2\tau^{-1}$$

then

$$\mathbb{E}1_{\mathcal{N}}\langle 1_A, p_{\omega X,\Lambda''}\rangle_{L^2(\mu)} > \mathbb{E}\alpha(1+\tau d/4)\int p_{X,\Lambda''}d\mu'.$$

By averaging there is some element such that the integrand on the left is at least that on the right and the result follows on setting $\Lambda''' = \{\lambda \in \Lambda'' : X_\lambda = 1\}$, and $\omega' = \omega|_{\Lambda'''}$. □

It is possible to optimise this lemma slightly more efficiently but we shall not concern ourselves with that here.

5. Basic properties of Bohr sets

In this section we collect the basic facts which we need about Bohr sets. Although this material has all now become standard this is in no small part due to the expository work of Tao; [TV06, Chapter 4.4] is the recommended reference.

To start with we have the following averaging argument *c.f.* [TV06, Lemma 4.20].

Lemma 5.1 (Size of Bohr sets)**.** *Suppose that B is a Bohr set. Then*

$$\mu_G(B_{2\rho}) = \exp(O(\mathrm{rk}(B)))\mu_G(B_\rho) \textit{ for all } \rho \in \mathbb{R}^+$$

and

$$\mu_G(B) \geqslant \exp(-O(\mathrm{rk}(B)))\prod_{\gamma\in\Gamma}\delta_\gamma.$$

Given two Bohr sets B and B' we define their *intersection* to be the Bohr set with frequency set $\Gamma \cup \Gamma'$ and width function $\delta \wedge \delta'$. It follows immediately from the previous lemma that Bohr sets are well behaved under intersections.

Lemma 5.2 (Intersections of Bohr sets)**.** *Suppose that B and B' are Bohr sets. Then*

$$\mathrm{rk}(B \wedge B') \leqslant \mathrm{rk}(B) + \mathrm{rk}(B')$$

and

$$\mu_G(B \wedge B') \geqslant \exp(-O(\mathrm{rk}(B) + \mathrm{rk}(B')))\mu_G(B)\mu_G(B').$$

Proof. The rank inequality is obvious. For the density note that

$$(B \wedge B')_1 \supset (B_{1/2} - B_{1/2}) \cap (B'_{1/2} - B'_{1/2})$$

by the triangle inequality. On the other hand

$$\mu_G((B_{1/2} - B_{1/2}) \cap (B'_{1/2} - B'_{1/2}))\mu_G(B_{1/2})\mu_G(B'_{1/2})$$

is at least

$$\int 1_{B_{1/2}} * 1_{B_{1/2}}1_{B'_{1/2}} * 1_{B'_{1/2}}d\mu_G = \|1_{B_{1/2}} * 1_{B'_{1/2}}\|^2_{L^2(\mu_G)} \geqslant (\mu_G(B_{1/2})\mu_G(B'_{1/2}))^2,$$

where the last inequality is Cauchy-Schwarz. The result now follows from the first case of the previous lemma. □

It is the first condition of Lemma 5.1 which is the most important and informs the definition of the dimension of a Bohr set: a Bohr set B is said to be *d-dimensional* if

$$\mu_G(B_{2\rho}) \leqslant 2^d \mu_G(B_\rho) \text{ for all } \rho \in (0,1],$$

and it will be convenient to assume that it is always at least 1. Since any d-dimensional Bohr set is always d'-dimensional for all $d' \geqslant d$ this is not a problem.

Note, in particular, that by Lemma 5.1 a Bohr set of rank k has dimension $O(k)$. It is, however, the dimension which is the important property of Bohr sets and the only reason we mention the rank is that it is sub-additive with respect to intersection, unlike dimension.

Some Bohr sets behave better than others: a d-dimensional Bohr set is said to be *C-regular* if

$$\frac{1}{1+Cd|\eta|} \leqslant \frac{\mu_G(B_{1+\eta})}{\mu_G(B_1)} \leqslant 1 + Cd|\eta| \text{ for all } \eta \text{ with } |\eta| \leqslant 1/Cd.$$

Crucially, regular Bohr sets are plentiful:

Lemma 5.3 (Regular Bohr sets)**.** *There is an absolute constant $C_{\mathcal{R}}$ such that whenever B is a Bohr set, there is some $\lambda \in [1/2, 1)$ such that B_λ is $C_{\mathcal{R}}$-regular.*

The result is proved by a covering argument due to Bourgain [Bou99]; for details one may also consult [TV06, Lemma 4.25]. For the remainder of the paper we shall say *regular* for $C_{\mathcal{R}}$-regular.

6. Bohr sets as majorants

One of the key properties of Bohr sets is that they can be used to majorise sets of characters. They are particularly good at this because if two sets are majorised by two different Bohr sets then their sumset is majorised by (a dilation of) the intersection of the Bohr sets.

We begin by recalling a lemma which shows that a set majorised by a Bohr set is approximately annihilated by a dilate, and if a set is approximately annihilated by a Bohr set then it is certainly majorised by it.

Lemma 6.1 (Majorising and annihilating)**.** *Suppose that B is a regular d-dimensional Bohr set. Then*

$$\{\gamma : |\widehat{\beta}(\gamma)| \geqslant \kappa\} \subset \{\gamma : |1-\gamma(x)| \leqslant C_{6.1}\rho\kappa^{-1}d \textit{ for all } x \in B_\rho\},$$

and

$$\{\gamma : |1-\gamma(x)| \leqslant \eta \textit{ for all } x \in B\} \subset \{\gamma : |\widehat{\beta}(\gamma)| \geqslant 1-\eta\}.$$

Proof. First, suppose that $|\widehat{\beta}(\gamma)| \geqslant \kappa$ and $y \in B_\rho$. Then

$$|1-\gamma(y)|\kappa \leqslant |\int \gamma(x)d\beta - \int \gamma(x+y)d\beta(x)| \leqslant \frac{\mu_G(B_{1+\rho} \setminus B_{1-\rho})}{\mu_G(B)} = O(d\rho)$$

provided $\rho \leqslant 1/C_{\mathcal{R}}d$. The first inclusion follows; the second is a trivial application of the triangle inequality. □

The proof is from [GK09, Lemma 3.6] where the significance of this result seems to have first been clearly recognised.

Of particular interest to us is the majorising of orthogonal and dissociated sets and in light of the previous lemma that can be achieved by the following. Recall that if X is a set then $\mathcal{P}(X)$ denotes the *power set*, that is the set of all subsets of X.

Lemma 6.2 (Annihilating orthogonal sets)**.** *Suppose that B is a regular d-dimensional Bohr set and Δ is a set of characters with (η,β)-relative size k. Then there is a set Λ, η-orthogonal w.r.t. β, and an order-preserving map $\phi:\mathcal{P}(\Lambda)\rightarrow\mathcal{P}(\Delta)$ such that $\phi(\Lambda)=\Delta$ and for any $\Lambda'\subset\Lambda$ and $\gamma\in\phi(\Lambda')$ we have*

$$|1-\gamma(x)|\leqslant C_{6.2}(\nu+\rho d^2\eta^{-1}\log 2k\eta^{-1}) \text{ for all } x\in B_\rho\wedge B'_\nu$$

where B' is the Bohr set with constant width function 2 and frequency set Λ'.

Proof. Let $L:=\lceil\log_2 2(k+1)^{3/2}\eta^{-1}\rceil$, the reason for which choice will become apparent, and define

$$\beta^+:=\beta_{1+L\rho'}*\beta_{\rho'}*\cdots*\beta_{\rho'},$$

where $\beta_{\rho'}$ occurs L times in the expression. By regularity (of B) we can pick $\rho'\in(\Omega(\eta/dL),1]$ such that $B_{\rho'}$ is regular and $\mu_G(B_{1+L\rho'})\leqslant(1+\eta/3)\mu_G(B)$. On the other hand we have the point-wise inequality

$$\mu_G(B)\beta=1_B\leqslant 1_{B_{1+L\rho'}}*\beta_{\rho'}*\cdots*\beta_{\rho'}=\mu_G(B_{1+L\rho'})\beta^+,$$

whence $\beta\leqslant 1+\eta/3)\beta^+$. It follows that if Λ is $\eta/2$-orthogonal w.r.t. β^+ then Λ is η-orthogonal w.r.t. β, and hence Λ has size at most k. From now on all orthogonality will be w.r.t. β^+.

We put $\eta_i:=i\eta/2(k+1)$ and begin by defining a sequence of sets $\Lambda_0,\Lambda_1,\dots$ iteratively such that Λ_i is η_i-orthogonal. We let $\Lambda_0:=\emptyset$ which is easily seen to be 0-orthogonal. Now, suppose that we have defined Λ_i as required. If there is some $\gamma\in\Delta\setminus\Lambda_i$ such that $\Lambda_i\cup\{\gamma\}$ is η_{i+1}-orthogonal then let $\Lambda_{i+1}:=\Lambda_i\cup\{\gamma\}$. Otherwise, terminate the iteration.

Note that for all $i\leqslant k+1$, if the set Λ_i is defined then it is certainly $\eta/2$-orthogonal and so $|\Lambda_i|\leqslant k$. However, if the iteration had continued for $k+1$ steps then $|\Lambda_{k+1}|>k$. This contradiction means that there is some $i\leqslant k$ such that $\Lambda:=\Lambda_i$ is η_i-orthogonal and $\Lambda_i\cup\{\gamma\}$ is not η_{i+1}-orthogonal for any $\gamma\in\Delta\setminus\Lambda_i$.

Now, we define ϕ in $\mathcal{P}(\Lambda)$ by

$$\phi(\Lambda')=\{\gamma\in\Delta:|\widehat{\beta_{\rho'}}(\gamma-\lambda)|\geqslant 1/2 \text{ for some } \lambda\in\Lambda'\cup\{0_{\widehat{G}}\}\}.$$

It is immediate that ϕ is order-preserving. Moreover we have the following claim.

Claim. $\phi(\Lambda)=\Delta$.

Proof. If $\lambda\in\Lambda$ then $\lambda\in\phi(\Lambda)$ since $\widehat{\beta_{\rho'}}(0_{\widehat{G}})=1$, so we need only consider $\gamma\in\Delta\setminus\Lambda$. In that case there is some $g\in\ell^2(\Lambda)$ and $\nu\in\mathbb{C}$ such that

$$\int|1+g^\vee+\nu\gamma|^2d\beta^+>(1+\eta_{i+1})(1+\|g\|_{\ell^2(\Lambda)}^2+|\nu|^2).$$

It follows from the η_i-orthogonality of Λ that

$$\begin{aligned} 2\operatorname{Re}\langle 1+g^{\vee},\nu\gamma\rangle_{L^2(\beta^+)} &> (1+\eta_{i+1})(1+\|g\|_{\ell^2(\Lambda)}^2+|\nu|^2)\\ &\quad -(1+\eta_i)(1+\|g\|_{\ell^2(\Lambda)}^2)-|\nu|^2\\ &\geqslant \frac{\eta}{2(k+1)}(1+\|g\|_{\ell^2(\Lambda)}^2+|\nu|^2)\\ &\geqslant \frac{\eta}{(k+1)}(1+\|g\|_{\ell^2(\Lambda)}^2)^{1/2}|\nu|. \end{aligned}$$

Thus

$$|\langle 1+g^{\vee},\gamma\rangle_{L^2(\beta^+)}|\geqslant \frac{\eta}{2(k+1)}(1+\|g\|_{\ell^2(\Lambda)}^2)^{1/2}$$

By Plancherel's theorem this tells us that

$$\begin{aligned} \frac{\eta}{2(k+1)}(1+\|g\|_{\ell^2(\Lambda)}^2)^{1/2} &\leqslant \left|\widehat{\beta^+}(\gamma)+\sum_{\lambda\in\Lambda}g(\lambda)\widehat{\beta^+}(\gamma-\lambda)\right|\\ &\leqslant (1+\|g\|_{\ell^1(\Lambda)})\sup_{\lambda\in\Lambda\cup\{0_{\widehat{G}}\}}|\widehat{\beta_{\rho'}}(\gamma-\lambda)|^L. \end{aligned}$$

By Cauchy-Schwarz we have $1+\|g\|_{\ell^1(\Lambda)}\leqslant\sqrt{k+1}(1+\|g\|_{\ell^2(\Lambda)}^2)^{1/2}$, and so it follows (given the size of L) that there is some $\lambda\in\Lambda\cup\{0_{\widehat{G}}\}$ such that $|\widehat{\beta_{\rho'}}(\gamma-\lambda)|\geqslant 1/2$. □

Now, suppose that $\Lambda'\subset\Lambda$ and $\gamma\in\phi(\Lambda')$. It follows that there is some $\lambda\in\Lambda'\cup\{0_{\widehat{G}}\}$ such that $|\widehat{\beta_{\rho'}}(\gamma-\lambda)|\geqslant 1/2$ and thus by Lemma 6.1 we have that

$$|1-\gamma(x)\overline{\lambda}(x)|\leqslant 2C_{6.1}\rho'^{-1}\rho d \text{ for all } x\in B_\rho.$$

On the other hand by definition $|1-\lambda(x)|=O(\nu)$ if $x\in B'_\nu$ and so the result follows from the triangle inequality. □

The argument for dissociated sets is rather similar so we omit many of the details.

Lemma 6.3 (Annihilating dissociated sets)**.** *Suppose that B is a regular d-dimensional Bohr set and Δ is a set of characters with (η,β)-relative entropy k. Then there is a set Λ, η-dissociated w.r.t. β, such that for all $\gamma\in\Delta$ we have*

$$|1-\gamma(x)|\leqslant C_{6.3}(k\nu+\rho d^2\eta^{-1}(k+\log 2\eta^{-1})) \text{ for all } x\in B_\rho\wedge B'_\nu$$

where B' is the Bohr set with constant width function 2 and frequency set Λ.

Proof. By the same argument as at the start of the previous lemma, but with $L=\lceil\log_2 3^k 2(k+1)\eta^{-1}\rceil$, we get some $i\leqslant k$, an η_i-dissociated (w.r.t. β^+) set Λ of at most k characters, such that for all $\gamma\in\Delta\setminus\Lambda$ there is a function $\omega:\Lambda\to D$ and $\nu\in D$ such that

$$\int p_{\omega,\Lambda}(1+\operatorname{Re}(\nu\gamma))d\beta^+>\exp(\eta_{i+1}).$$

Now, suppose that $\gamma \in \Delta$. If $\gamma \in \Lambda$ then the conclusion is immediate, so we may assume that $\gamma \in \Delta \setminus \Lambda$. Then, since Λ is η_i-dissociated, we see that

$$|\int p_{\omega,\Lambda}\overline{\gamma}d\beta^+| > \exp(\eta_{i+1}) - \exp(\eta_i) \geqslant \frac{\eta}{2(k+1)}.$$

Applying Plancherel's theorem we get that

$$\frac{\eta}{2(k+1)} \leqslant \left|\sum_{\lambda \in \mathrm{Span}(\Lambda)} \widehat{p_{\omega,\Lambda}}(\lambda)\widehat{\beta^+}(\gamma - \lambda)\right| \leqslant 3^k \sup_{\lambda \in \mathrm{Span}(\Lambda)} |\widehat{\beta}_\rho(\gamma-\lambda)|^L.$$

Given the choice of L there is some $\lambda \in \mathrm{Span}(\Lambda)$ such that $|\widehat{\beta_{\rho'}}(\gamma - \lambda)| \geqslant 1/2$. The argument now proceeds as before on noting that $|1 - \lambda(x)| = O(k\nu)$ if $x \in B'_\nu$ by the triangle inequality. □

It is also useful to record an immediate corollary of this.

Corollary 6.4. *Suppose that B is a regular d-dimensional Bohr set and Δ is a set of characters with (η, β)-relative entropy k. Then there is a Bohr set $B' \leqslant B$ with*

$$\mathrm{rk}(B') \leqslant \mathrm{rk}(B) + k \text{ and } \mu_G(B') \geqslant (\eta/2dk)^{O(d)}(1/2k)^{O(k)}\mu_G(B)$$

such that $|1 - \gamma(x)| \leqslant 1/2$ for all $x \in B'$ and $\gamma \in \Delta$.

7. Getting a density increment

The basic dichotomy of the density increment strategy is driven by the following lemma which says that either we are in a situation where we have many three-term progressions or else we have a concentration of Fourier mass.

Lemma 7.1. *Suppose that B is a regular d-dimensional Bohr set and $A \subset B$ and $A' \subset B_\rho$ have density $\alpha > 0$ and $\alpha' > 0$ respectively. Then either*

(i) *(Many three-term progressions)*

$$T(1_A, 1_{A'}, 1_A) \geqslant \alpha^2\alpha'\mu_G(B)\mu_G(B_\rho)/2;$$

(ii) *(Concentration of Fourier mass)*

$$\sum_{-2\gamma \in \mathrm{Spec}_{c_{7.1}\alpha}(1_{A'},\beta_\rho)} |((1_A - \alpha)1_B)^\wedge(\gamma)|^2|(1_{A'}d\beta_\rho)^\wedge(-2\gamma)| \geqslant c_{7.1}\alpha^2\alpha'\mu_G(B);$$

provided $\rho \leqslant c'_{7.1}\alpha/d$.

Proof. Write $M : L^2(\beta) \to L^2(\beta)$ for the linear operator taking $f \in L^2(\beta)$ to

$$x \mapsto \int f(x - 2y)1_{A'}(y)d\beta_\rho(y)1_B(x),$$

and put

$$T(f) := \langle Mf, f\rangle_{L^2(\beta)},$$

so that

$$T(1_A, 1_{A'}, 1_A) = \mu_G(B)\mu_G(B_\rho)T(1_A).$$

We begin by noting that

$$\begin{aligned} T(1_A) &= T((1_A - \alpha)1_B) + \alpha\langle M1_A, 1_B\rangle_{L^2(\beta)} \\ &\quad +\alpha\langle M1_B, 1_A\rangle_{L^2(\beta)} - \alpha^2\langle M1_B, 1_B\rangle_{L^2(\beta)}. \end{aligned}$$

By regularity we have

$$\|M1_B - \alpha' 1_B\| = O(\rho d\alpha'\mu_G(B))$$

provided $\rho \leqslant 1/2C_{\mathcal{R}}d$. Thus, by the triangle inequality we have

$$T(1_A) = T((1_A - \alpha)1_B) + \alpha^2\alpha' + O(\alpha\alpha'\rho d).$$

It follows that there is some absolute $c > 0$ such that if $\rho \leqslant c\alpha/d$ then either we are in the first case of the lemma or else

$$|T((1_A - \alpha)1_B)| \geqslant \alpha^2\alpha'/4;$$

we may naturally assume the latter or else we are done. The usual application of the transform tells us that

$$\sum_{\gamma\in\widehat{G}} |((1_A - \alpha)1_B)^\wedge(\gamma)|^2|(1_{A'}d\beta_\rho)^\wedge(-2\gamma)| \geqslant \alpha^2\alpha'\mu_G(B)/4,$$

from which it follows that we are in the second case of the lemma by the triangle inequality. □

The remaining two lemmas of the section encode the other half of the density increment strategy: having analysed the Fourier mass in some way, we need results which convert this analysis into a density increment.

First we have the Heath-Brown-Szemerédi energy increment technique. There are more general versions of the lemma suitable for other problems; the one here is adapted to our circumstance and, in particular, is the only place we use the fact that G has odd order.

Lemma 7.2 (Energy to density lemma)**.** *Suppose that B is a regular d-dimensional Bohr set, $A \subset B$ has density $\alpha > 0$, $\mathcal{L}$ is a set of characters such that*

$$\sum_{-2\gamma\in\mathcal{L}} |((1_A - \alpha)1_B)^\wedge(\gamma)|^2 \geqslant K\alpha^2\mu_G(B),$$

and $B' \leqslant B_\rho$ is a d'-dimensional Bohr set such that

$$\mathcal{L} \subset \{\gamma : |1 - \gamma(x)| \leqslant 1/2 \text{ for all } x \in B'\}.$$

Then there is a regular Bohr set B'' with

$$\operatorname{rk}(B'') = \operatorname{rk}(B') \text{ and } \mu_G(B'') \geqslant 2^{-O(d')}\mu_G(B')$$

such that

$$\|1_A * \beta''\|_{L^\infty(\mu_G)} \geqslant \alpha(1 + c_{7.2}K)$$

provided $\rho \leqslant c'_{7.2}\alpha K/d$.

Proof. Put $B''' := -2.B'_{1/2}$, and note that B''' has the same rank and size as $B'_{1/2}$ (since G has odd order), and

$$\{\gamma : -2\gamma \in \mathcal{L}\} \subset \{\gamma : |1-\gamma(-2x)| \leqslant 1/2 \text{ for all } x \in B'\}.$$

Pick $\nu \in [1/2,1)$ such that $B'' := B'''_\nu$ is regular and so, since $\{-2x : x \in B'\} \supset \{x : x \in B''\}$, we conclude that

$$\{\gamma : -2\gamma \in \mathcal{L}\} \subset \{\gamma : |1-\gamma(x)| \leqslant 1/2 \text{ for all } x \in B''\}.$$

By the triangle inequality if γ is in this second set, then $|\widehat{\beta''}(\gamma)| \geqslant 1/2$. Thus Plancherel's theorem tells us that

$$\begin{aligned}\|((1_A-\alpha)1_B)*\beta''\|^2_{L^2(\mu_G)} &= \sum_{\gamma\in\widehat{G}} |((1_A-\alpha)1_B)^\wedge(\gamma)|^2|\widehat{\beta''}(\gamma)|^2\\ &\geqslant \frac{1}{4}\sum_{-2\gamma\in\mathcal{L}} |((1_A-\alpha)1_B)^\wedge(\gamma)|^2 \geqslant \frac{K}{4}\alpha^2\mu_G(B).\end{aligned}$$

But $B'' \subset B'_1 \subset B_\rho$, so by regularity we have that

$$\langle f, 1_B*\beta''*\beta''\rangle_{L^2(\mu_G)} = (\int f d\beta + O(\rho d\|f\|_{L^\infty(\beta)}))\mu_G(B)$$

for any $f \in L^\infty(\beta)$ provided $\rho \leqslant 1/2C_{\mathcal{R}}d$. Thus

$$\begin{aligned}\|((1_A-\alpha)1_B)*\beta''\|^2_{L^2(\mu_G)} &= \|1_A*\beta''\|^2_{L^2(\mu_G)} - 2\alpha((\alpha+O(\rho d))\mu_G(B)\\ &\quad +\alpha^2(1+O(\rho d))\mu_G(B),\end{aligned}$$

whence

$$\|1_A*\beta''\|^2_{L^2(\mu_G)} \geqslant \alpha^2\mu_G(B) + K\alpha^2\mu_G(B)/4 - O(\alpha\rho d\mu_G(B)).$$

The result now follows provided ρ is sufficiently small on applying Hölder's inequality and dividing by $\alpha\mu_G(B)$. □

The second lemma we need is used for leveraging of the second conclusion of Lemma 4.7 and gives a way of passing between correlation with a Riesz product and a density increment on a Bohr set.

Lemma 7.3 (Riesz product correlation to density lemma)**.** *Suppose that B is a regular d-dimensional Bohr set, $A \subset B_{1-3\rho}$, Λ is a set of k characters and $\omega : \Lambda \to D$ is such that*

$$\langle 1_A, p_{\omega,\Lambda}\rangle_{L^2(\beta)} \geqslant \alpha(1+\epsilon)\int p_{1,\Lambda} d\mu,$$

where $\mu = \beta_\rho * \beta_\rho$. Then there is a regular Bohr set B' with

$$\mathrm{rk}(B') \leqslant \mathrm{rk}(B) + k \text{ and } \mu_G(B') \geqslant (\epsilon\alpha/2k)^{O(k)}2^{-O(d)}\mu_G(B_\rho)$$

such that $\|1_A*\beta'\|_{L^\infty(\mu_G)} \geqslant \alpha(1+\epsilon/2)$.

Proof. Let B'' be the Bohr set with frequency set Λ and width function equal to the constant function 1. Let $\lambda_1, \dots, \lambda_k$ be an enumeration of Λ and put

$$\omega_i : \Lambda \to D; \lambda \mapsto \begin{cases} \omega_\lambda & \text{if } \lambda = \lambda_j \text{ for some } j > i \\ 0 & \text{if } \lambda = \lambda_i \\ \omega_\lambda \lambda(y) & \text{if } \lambda = \lambda_j \text{ for some } j < i. \end{cases}$$

With these definitions we have

$$p_{\omega,\Lambda}(x+y) - p_{\omega,\Lambda}(x) = \sum_{i=1}^{k} (\mathrm{Re}(\omega_{\lambda_i}\lambda_i(x+y)) - \mathrm{Re}(\omega_{\lambda_i}\lambda_i(x)))p_{\omega_i,\Lambda}(x),$$

but

$$|\,\mathrm{Re}(\omega_{\lambda_i}\lambda_i(x+y)) - \mathrm{Re}(\omega_{\lambda_i}\lambda_i(x))| = O(\rho')$$

if $y \in B''_{\rho'}$. Since μ is positive definite we see that Λ is K-dissociated w.r.t. μ with $K = \log \int p_{1,\Lambda} d\mu$. Thus, integrating the above, we get that

$$\int |\tau_y(p_{\omega,\Lambda}) - p_{\omega,\Lambda}| d\beta * \mu \leqslant \sum_{i=1}^{k} O(\rho'). \int p_{\omega_i,\Lambda} d\beta * \mu = O(\rho' k \exp(K))$$

for all $y \in B''_{\rho'}$ by monotonicity of dissociativity. Let $\rho' = \Omega(\epsilon\alpha/k)$ be such that the above error term is at most $\epsilon\alpha\exp(K)/2$ and put $B' = (B''_{\rho'} \wedge B_\rho)_\nu$ where $\nu \in [1/2, 1)$ is such that B' is regular. It follows by the triangle inequality that

$$\int |p_{\omega,\Lambda} * \beta' - p_{\omega,\Lambda}| d\beta * \mu \leqslant \epsilon\alpha\exp(K)/2.$$

On the other hand, since $A \subset B_{1-2\rho}$ we have that

$$\langle 1_A, p_{\omega,\Lambda}\rangle_{L^2(\beta)} = \langle 1_A, p_{\omega,\Lambda}\rangle_{L^2(\beta*\mu)},$$

and so by the triangle inequality and hypothesis we have

$$\langle 1_A, p_{\omega,\Lambda} * \beta'\rangle_{L^2(\beta)} \geqslant \alpha(1+\epsilon/2)\exp(K).$$

However, since $A \subset B_{1-3\rho}$, we have that

$$\langle 1_A, p_{\omega,\Lambda} * \beta'\rangle_{L^2(\beta)} = \langle 1_A * \beta', p_{\omega,\Lambda}\rangle_{L^2(\beta*\mu)}.$$

Finally, by monotonicity of dissociativity again we get that

$$\|1_A * \beta'\|_{L^\infty(\mu_G)} \exp(K) \geqslant \alpha(1+\epsilon/2)\exp(K),$$

and the result follows. □

8. Roth's theorem in high rank Bohr sets

From now on we assume that the underlying group has odd order so that we may apply Lemma 7.2.

The purpose of this section is to establish the following theorem which is the formal version of the result discussed in §3.2.

Theorem 8.1. *Suppose that B is a rank k Bohr set and $A \subset B$ has density $\alpha > 0$. Then*

$$T(1_A, 1_A, 1_A) = \exp(-\widetilde{O}(k + \alpha^{-2}))\mu_G(B)^2.$$

We shall prove this iteratively using the following lemma.

Lemma 8.2. *Suppose that B is a d-dimensional Bohr set such that B and B_ρ are regular, and $A \subset B$ and $A' \subset B_\rho$ have density $\alpha > 0$ and $\alpha' > 0$ respectively with*[3] *$\alpha = \Theta(\alpha')$. Then at least one of the following is true:*

(i) *(Many three-term progressions)*

$$T(1_A, 1_{A'}, 1_A) \geqslant \alpha^2\alpha'\mu_G(B)\mu_G(B_\rho)/2;$$

(ii) *(Energy induced density increment) or there is a regular Bohr set B' with*

$$\mathrm{rk}(B') \leqslant \mathrm{rk}(B) + \widetilde{O}(\alpha^{-2}) \text{ and } \mu_G(B') \geqslant \exp(-\widetilde{O}(d + \alpha^{-2}))\mu_G(B_\rho)$$

such that $\|1_A * \beta'\|_{L^\infty(\mu_G)} \geqslant \alpha(1 + c_{8.2})$;

provided $\rho \leqslant c'_{8.2}\alpha/d$.

Proof. By Lemma 7.1 (provided $\rho \leqslant c'_{7.1}\alpha/d$) we are either in the first case of the lemma or else

$$\sum_{-2\gamma \in \mathrm{Spec}_{c_{7.1}\alpha}(1_{A'}, \beta_\rho)} |((1_A - \alpha)1_B)^\wedge(\gamma)|^2 \geqslant c_{7.1}\alpha^2\mu_G(B). \tag{8.1}$$

By (the relative version of) the Chang bound (Lemma 4.6) we see that $\mathrm{Spec}_{c_{7.1}\alpha}(1_{A'}, \beta_\rho)$ has $(1, \beta_\rho)$-relative entropy $O(\alpha^{-2}\log\alpha'^{-1}) = \widetilde{O}(\alpha^{-2})$. Thus, by Corollary 6.4 there is a Bohr set $B'' \leqslant B_\rho$ with

$$\mathrm{rk}(B'') \leqslant \mathrm{rk}(B) + \widetilde{O}(\alpha^{-2}) \text{ and } \mu_G(B'') \geqslant \exp(-\widetilde{O}(d + \alpha^{-2}))\mu_G(B_\rho)$$

such that

$$|1 - \gamma(x)| \leqslant 1/2 \text{ for all } x \in B'' \text{ and } \gamma \in \mathrm{Spec}_{c_{7.1}\alpha}(1_{A'}, \beta_\rho).$$

It follows by Lemma 7.2 (with $\mathcal{L} := \mathrm{Spec}_{c_{7.1}\alpha}(1_{A'}, \beta_\rho)$ provided $\rho \leqslant c'_{7.2}c_{7.1}\alpha/d$) that we are in the second case of the lemma. □

---

[3] This condition is unnecessary but it makes for simpler statements, particularly later on. It is also a slight abuse; to be formally correct the reader may wish to take, for example, $2\alpha' \geqslant \alpha \geqslant \alpha'/2$.

Proof of Theorem 8.1. We proceed iteratively to construct a sequence of regular Bohr sets $B^{(i)}$ of rank k_i and dimension $d_i = O(k_i)$. We write

$$\alpha_i := \|1_A * \beta^{(i)}\|_{L^\infty(\mu_G)},$$

put $B^{(0)} := B$ and suppose that we have defined $B^{(i)}$. By regularity we have that

$$(1_A * \beta^{(i)}_{\rho_i} + 1_A * \beta^{(i)}_{\rho_i\rho_i'}) * \beta^{(i)}(x_i) = 2\alpha_i + O(\rho_i k_i)$$

for some x_i. Pick $\rho_i = \Omega(\alpha_i/k_i)$ and then $\rho_i' = \Omega(\alpha_i/k_i)$ such that $B^{(i)}_{\rho_i}$ and $B^{(i)}_{\rho_i'\rho_i}$ are regular,

$$(1_A * \beta^{(i)}_{\rho} + 1_A * \beta^{(i)}_{\rho_i\rho_i'} * \beta^{(i)}(x_i) \geqslant 2\alpha_i(1 - c_{8.2}/4),$$

and $\rho_i' \leqslant c_{8.2}'\alpha_i/d_i$. Thus there is some x_i' such that

$$1_A * \beta^{(i)}_{\rho}(x_i') + 1_A * \beta^{(i)}_{\rho_i\rho_i'}(x_i') \geqslant 2\alpha_i(1 - c_{8.2}/4).$$

If one of the two terms is at least $\alpha_i(1 + c_{8.2}/4))$ then set $B^{(i+1)}$ to be $B^{(i)}_{\rho_i}$ or $B^{(i)}_{\rho_i'\rho_i}$ respectively. Otherwise we see that $A_i := (x_i' - A) \cap B^{(i)}_{\rho_i}$ has

$$\beta^{(i)}_{\rho_i}(A_i) \geqslant \alpha_i(1 - 3c_{8.2}/4),$$

and $A_i' := (x_i' - A) \cap B^{(i)}_{\rho_i'\rho_i}$ has

$$\beta^{(i)}_{\rho_i'\rho_i}(A_i') \geqslant \alpha_i(1 - 3c_{8.2}/4).$$

By Lemma 8.2 we see that either we are done in that

$$T(1_A, 1_A, 1_A) = \Omega(\alpha_i^3 \mu_G(B^{(i)}_{\rho_i})\mu_G(B^{(i)}_{\rho_i'\rho_i})) \tag{8.2}$$

or there is a regular Bohr set $B^{(i+1)}$ such that

$$k_{i+1} \leqslant k_i + \widetilde{O}(\alpha_i^{-2}),$$

$$\mu_G(B^{(i+1)}) \geqslant \exp(-\widetilde{O}(k_i + \alpha_i^{-2}))\mu_G(B^{(i)}_{\rho_i'\rho_i}) \geqslant \exp(-\widetilde{O}(k_i + \alpha_i^{-2}))\mu_G(B^{(i)}),$$

and $\alpha_{i+1} \geqslant \alpha_i(1 + c_{8.2}/4)$. Thus, if the iteration proceeds to the next step, then however it did so we have $\alpha_{i+1} \geqslant \alpha_i(1 + \Omega(1))$. This cannot happen for more than $O(\log 2\alpha^{-1})$ step, whence we are in (8.2) at some step $i_0 = O(\log 2\alpha^{-1})$, where we shall have

$$k_{i_0} = k + \widetilde{O}(\alpha^{-2}) \text{ and } \mu_G(B^{(i_0)}) \geqslant \exp(-\widetilde{O}(k + \alpha^{-2}))\mu_G(B),$$

from which the result follows. □

9. The main result

In this section we prove the following theorem which is the main result of the paper.

Theorem 9.1. *Suppose that A has density $\alpha > 0$. Then*

$$T(1_A, 1_A, 1_A) = \exp(-\widetilde{O}(\alpha^{-4/3})).$$

To see how Theorem 1.1 follows simply note that if $A \subset \{1, \dots, N\}$ has density α and no non-trivial three-term progressions then it can be embedded in $\mathbb{Z}/(4N+1)\mathbb{Z}$ with density $\Omega(\alpha)$ such that it also has no non-trivial three-term progressions there, but then

$$(4N+1) \geqslant (4N+1)^2 \exp(-\widetilde{O}(\alpha^{-4/3})),$$

which can be rearranged to give the desired bound.

As before we proceed iteratively; the following is the driver.

Lemma 9.2. *Suppose that B is a d-dimensional Bohr set such that B and B_ρ are regular, and $A \subset B$ and $A' \subset B_\rho$ have density $\alpha > 0$ and $\alpha' > 0$ respectively with $\alpha = \Theta(\alpha')$. Then at least one of the following is true:*

(i) *(Many three-term progressions)*

$$T(1_A, 1_{A'}, 1_A) \geqslant \alpha^2 \alpha' \mu_G(B) \mu_G(B_\rho)/2;$$

(ii) *(Energy induced density increment from large coefficients) or there is a regular Bohr set B' with*

$$\mathrm{rk}(B') \leqslant \mathrm{rk}(B) + \widetilde{O}(\alpha^{-1}) \text{ and } \mu_G(B') \geqslant \exp(-\widetilde{O}(d + \alpha^{-1}))\mu_G(B_\rho)$$

such that $\|1_A * \beta'\|_{L^\infty(\mu_G)} \geqslant \alpha(1 + c_{9.2})$;

(iii) *(Terminal density increment) or there is a regular Bohr set B' with*

$$\mathrm{rk}(B') \leqslant \mathrm{rk}(B) + \widetilde{O}(\alpha^{-4/3}) \text{ and } \mu_G(B') \geqslant \exp(-\widetilde{O}(d + \alpha^{-4/3}))\mu_G(B_\rho)$$

such that $\|1_A * \beta'\|_{L^\infty(\mu_G)} = \widetilde{\Omega}(\alpha^{2/3})$;

(iv) *(Density increment from many independent small coefficients) or there is a regular Bohr set B' with*

$$\mathrm{rk}(B') \leqslant \mathrm{rk}(B) + \widetilde{O}(\alpha^{-2/3}) \text{ and } \mu_G(B') \geqslant \exp(-\widetilde{O}(d + \alpha^{-2/3}))\mu_G(B_\rho)$$

such that $\|1_{A'} * \beta'\|_{L^\infty(\mu_G)} \geqslant \alpha'(1 + c_{9.2}\alpha^{1/3})$;

provided $\rho \leqslant c'_{9.2}\alpha/d$.

Proof. The argument is not difficult although it involves a number of somewhat technical details. To start with (for the purpose of applying Lemma 7.3 later) we shall adjust the set A' slightly. Let $\rho' = c\alpha'^2/d$ for some absolute $c > 0$ such that $B_{\rho\rho'}$ is regular and

$$\mu_G(B_\rho \setminus B_{\rho(1-3\rho')}) \leqslant \alpha'^2/128$$

(possible by regularity of B_ρ), and put $A'' := A' \cap B_{\rho(1-3\rho')}$ so that

$$A'' \subset B_{\rho(1-3\rho')} \text{ and } \beta_\rho(A'') \geqslant \alpha'(1 - \alpha'/128).$$

By Lemma 7.1 applied to A and A'' (provided $\rho \leqslant c'_{7.1}\alpha/d$) we are either in the first case of the lemma or else

$$\sum_{-2\gamma\in\mathrm{Spec}_{c_{7.1}\alpha}(1_{A''},\beta_\rho)} |((1_A-\alpha)1_B)^\wedge(\gamma)|^2|(1_{A''}d\beta_\rho)^\wedge(-2\gamma)| \geqslant c_{7.1}\alpha^2\alpha'\mu_G(B). \tag{9.1}$$

We assume that we are not in the first case and proceed with analysis of $\mathrm{Spec}_{c_{7.1}\alpha}(1_{A''},\beta_\rho)$ depending on where the mass in (9.1) concentrates. Let $\epsilon := c\alpha^{1/2}$ where c is an absolute constant such that

$$C_{4.7}\log 2c_{7.1}^{-1}\tau^{-1} \leqslant c_{4.7}\tau^{-2/3} \text{ whenever } \tau \leqslant c. \tag{9.2}$$

First we split into two cases.

Case. *At least half the mass in (9.1) is supported by those* γ *with* $-2\gamma \in \mathrm{Spec}_\epsilon(1_{A''},\beta_\rho)$.

Analysis of case. We proceed as in the proof of Lemma 8.2: by (the relative version of) the Chang bound (Lemma 4.6) we see $\mathrm{Spec}_\epsilon(1_{A''},\beta_\rho)$ has $(1,\beta_\rho)$-relative entropy $\widetilde{O}(\alpha^{-1})$. Thus, by Corollary 6.4 there is a Bohr set $B'' \leqslant B_\rho$ with

$$\mathrm{rk}(B'') \leqslant \mathrm{rk}(B) + \widetilde{O}(\alpha^{-1}) \text{ and } \mu_G(B'') \geqslant \exp(-\widetilde{O}(d+\alpha^{-1}))\mu_G(B_\rho)$$

such that

$$|1-\gamma(x)| \leqslant 1/2 \text{ for all } x \in B'' \text{ and } \gamma \in \mathrm{Spec}_\epsilon(1_{A''},\beta_\rho).$$

It follows by Lemma 7.2 (provided $\rho \leqslant c'_{7.2}c_{7.1}\alpha/2d$) that we are in the second case of the lemma. □

Case. *At least half the mass in (9.1) is supported by those* γ *with* $-2\gamma \notin \mathrm{Spec}_\epsilon(1_{A''},\beta_\rho)$.

Analysis of case. By averaging there is some $\tau \in [c_{7.1}\alpha,\epsilon]$ such that the set of γ such that -2γ is a member of

$$\mathcal{V} := \{\gamma : 2\tau\alpha' > |(1_{A''}d\beta_\rho)^\wedge(\gamma)| \geqslant \tau\alpha'\}$$

supports at least $\Omega(\log^{-1} 2\alpha^{-1})$ of the mass in (9.1), whence

$$\sum_{-2\gamma\in\mathcal{V}} |((1_A-\alpha)1_B)^\wedge(\gamma)|^2 = \widetilde{\Omega}(\tau^{-1}\alpha^2)\mu_G(B).$$

Note that by definition of A'' we have that

$$|(1_{A''}d(\beta_\rho * \beta_{\rho\rho'}))^\wedge(\gamma)| = |(1_{A''}d\beta_\rho)^\wedge(\gamma)|,$$

whence $\mathcal{V} \subset \mathrm{Spec}_\tau(1_{A''},\beta_\rho * \beta_{\rho\rho'})$. However, the Bessel bound (Lemma 4.2) tells us that $\mathrm{Spec}_\tau(1_{A''},\beta_\rho * \beta_{\rho\rho'})$ has $(1,\beta_\rho * \beta_{\rho\rho'})$-relative size $O(\alpha^{-3})$ and so by monotonicity $\mathcal{V}$ has $(c'_{4.7}c_{7.1}\alpha,\beta_{\rho\rho'})$-relative size $O(\alpha^{-3})$.

Now, apply Lemma 6.2 to $\mathcal{V}$ and the Bohr set $B_{\rho\rho'}$ to get a set Λ of k characters, $(c'_{4.7}c_{7.1}\alpha)$-orthogonal w.r.t. $\beta_{\rho\rho'}$ (and hence $\mu := \beta_{\rho\rho'} * \beta_{\rho\rho'}$), and an order-preserving map $\phi : \mathcal{P}(\Lambda) \to \mathcal{P}(\mathcal{V})$ such that $\phi(\Lambda) = \mathcal{V}$ and for all $\Lambda' \subset \Lambda$ and $\gamma \in \phi(\Lambda')$ we have

$$|1-\gamma(x)| \leqslant 1/2 \text{ for all } x \in B_{\rho\rho'\rho''} \wedge B'''_\nu \tag{9.3}$$

where B''' is the Bohr set with constant width function 2 and frequency set Λ', and $\rho'' = \Omega(\alpha/d^2\log 2\alpha^{-1})$ and $\nu = \Omega(1)$.

We define a sub-additive weight w on $\mathcal{P}(\Lambda)$ by

$$w(\Lambda') = \sum_{-2\gamma\in\phi(\Lambda')} |((1_A-\alpha)1_B)^\wedge(\gamma)|^2,$$

and a sequence of sets $\Lambda_0,\dots,\Lambda_r$ as follows. We begin with $\Lambda_0 := \Lambda$, and at stage i if there is some $\Lambda'\subset\Lambda_i$ with

$$|\Lambda'| \geqslant |\Lambda_i|/2 \text{ and } w(\Lambda') \leqslant w(\Lambda)/\log 2k$$

then we put $\Lambda_{i+1} := \Lambda_i\setminus\Lambda'$. Since $|\Lambda_{i+1}| \leqslant |\Lambda_i|/2$ we see that the iteration terminates after at most $\log_2 k$ steps. At this point we have a set Λ_i such that every subset of size at least $|\Lambda_i|/2$ has weight at least $w(\Lambda_i)/\log 2k$. On the other hand

$$w(\Lambda_i) \geqslant \left(1-\frac{1}{\log 2k}\right)^{\lceil \log_2 k\rceil} w(\Lambda) = \widetilde{\Omega}(\tau^{-1}\alpha^2)\mu_G(B).$$

Once again we split into two cases.

Case. Λ_i *is small:* $|\Lambda_i| \leqslant \tau^{-1}\alpha^{-2/3}$.

Analysis of case. By averaging there is a subset $\Lambda'\subset\Lambda_i$ of size at most α^{-1} such that

$$w(\Lambda') \geqslant \frac{\lfloor\alpha^{-1}\rfloor}{|\Lambda_i|} w(\Lambda_i) = \widetilde{\Omega}(\alpha^{5/3})\mu_G(B).$$

It follows from (9.3) that there is a Bohr set $B'' \leqslant B$ with

$$\operatorname{rk}(B'') \leqslant \operatorname{rk}(B) + \alpha^{-1} \text{ and } \mu_G(B'') \geqslant \exp(-\widetilde{O}(d+\alpha^{-1}))\mu_G(B_\rho)$$

such that

$$|1-\gamma(x)| \leqslant 1/2 \text{ for all } x\in B'' \text{ and } \gamma\in\phi(\Lambda').$$

We apply Lemma 7.2 (provided $\rho \leqslant c'\alpha/d$ for some absolute c') and are certainly in the third case of the lemma. □

Case. Λ_i *is large:* $|\Lambda_i| \geqslant \tau^{-1}\alpha^{-2/3}$.

Analysis of case. We apply Lemma 4.7 with the set A'', measures β_ρ and μ, and parameter $m = \alpha^{-4/3}$ to get two possibilites.

Case. *There is a set* $\Lambda'\subset\Lambda_i$ *with* $|\Lambda'| \geqslant |\Lambda_i|/2$ *such that* Λ' *has* $(1,\mu)$*-relative entropy* $\widetilde{O}(\alpha^{-4/3})$.

Analysis of case. By monotonicity Λ' has $(1,\beta_{\rho\rho'})$-relative entropy $\widetilde{O}(\alpha^{-4/3})$, and so by Lemma 6.3 applied to Λ' and the Bohr set $B_{\rho\rho'}$ that there is a $\rho''' = (\alpha/2d)^{O(1)}$ and $\nu' = \Omega(1/k)$ such that

$$|1-\gamma(x)| \leqslant \nu'' \text{ for all } x\in (B_{\rho\rho'\rho'''}\wedge B''''_{\nu'})''_{\nu} \text{ and } \gamma\in\Lambda'$$

where B'''' is a Bohr set with constant width function 2 and a frequency set of size $\widetilde{O}(\alpha^{-4/3})$. It follows that $B''' \leqslant B_{\rho\rho'\rho'''} \wedge B''''_{\nu'}$, and inserting this in (9.3) we get that for all $\gamma \in \phi(\Lambda')$,

$$|1-\gamma(x)| \leqslant 1/2 \text{ for all } x \in B_{\rho\rho' \min\{\rho'',\rho'''\nu\}} \wedge B_{\nu'\nu}.$$

We conclude that there is a Bohr set $B'' \leqslant B_\rho$ with

$$\mathrm{rk}(B'') \leqslant \mathrm{rk}(B) + \widetilde{O}(\alpha^{-4/3}) \text{ and } \mu_G(B'') \geqslant \exp(-\widetilde{O}(d+\alpha^{-4/3}))\mu_G(B_\rho)$$

such that $|1-\gamma(x)| \leqslant 1/2$ for all $x \in B''$ and $\gamma \in \phi(\Lambda')$. On the other hand

$$w(\Lambda') \geqslant \frac{w(\Lambda_i)}{\log 2k} = \widetilde{\Omega}(\tau^{-1}\alpha^2)\mu_G(B) = \widetilde{\Omega}(\alpha^{3/2})\mu_G(B).$$

It follows from Lemma 7.2 that we are in the third case of the lemma provided $\rho \leqslant c''\alpha/d$ for some absolute c''. □

Case. *There is a set of characters Λ''' of $O(\sqrt{m})$ characters and a function $\omega' : \Lambda''' \to D$ such that*

$$\langle 1_{A''}, p_{\omega',\Lambda'''}\rangle_{L^2(\beta_\rho)} \geqslant \alpha'(1+\tau\sqrt{m}/4)\int p_{1,\Lambda'''}d\mu.$$

Analysis of case. Apply Lemma 7.3 and we are in the final case of the lemma. □

□

□

The proof is complete. □

Proof of Theorem 9.1. We proceed as before to construct a sequence of regular Bohr sets $B^{(i)}$ of rank k_i and dimension $d_i = O(k_i)$. We write

$$\alpha_i := \|1_A * \beta^{(i)}\|_{L^\infty(\mu_G)},$$

put $B^{(0)} := G$ and suppose that we have defined $B^{(i)}$. By regularity we have that

$$(1_A * \beta^{(i)}_{\rho_i} + 1_A * \beta^{(i)}_{\rho_i\rho_i'}) * \beta^{(i)}(x_i) = 2\alpha_i + O(\rho_i k_i)$$

for some x_i. Pick $\rho_i = \Omega(\alpha_i^{4/3}/k_i)$ and then $\rho_i' = \Omega(\alpha_i^{4/3}/k_i)$ such that $B^{(i)}_{\rho_i}$ and $B^{(i)}_{\rho_i'\rho_i}$ are regular,

$$(1_A * \beta^{(i)}_{\rho} + 1_A * \beta^{(i)}_{\rho_i\rho_i'}) * \beta^{(i)}(x_i) \geqslant 2\alpha_i(1 - c_{9.2}\alpha_i^{1/3}/4),$$

and $\rho_i' \leqslant c'_{9.2}\alpha_i/d_i$. Thus there is some x_i' such that

$$1_A * \beta^{(i)}_{\rho}(x_i') + 1_A * \beta^{(i)}_{\rho_i\rho_i'}(x_i') \geqslant 2\alpha_i(1 - c_{9.2}\alpha_i^{1/3}/4).$$

If one of the two terms is at least $\alpha_i(1+c_{9.2}\alpha_i^{1/3}/4))$ then set $B^{(i+1)}$ to be $B^{(i)}_{\rho_i}$ or $B^{(i)}_{\rho_i'\rho_i}$ respectively. Otherwise we see that $A_i := (x_i' - A) \cap B^{(i)}_{\rho_i}$ has

$$\beta^{(i)}_{\rho_i}(A_i) \geqslant \alpha_i(1 - 3c_{9.2}\alpha_i^{1/3}/4),$$

and $A_i' := (x_i' - A) \cap B^{(i)}_{\rho_i'\rho_i}$ has

$$\beta^{(i)}_{\rho_i'\rho_i}(A_i') \geqslant \alpha_i(1 - 3c_{9.2}\alpha_i^{1/3}/4).$$

By Lemma 9.2 we see that at least one of the following occurs:

(i) we stop the iteration with

$$T(1_A, 1_A, 1_A) = \Omega(\alpha_i^3 \mu_G(B^{(i)}_{\rho_i})\mu_G(B^{(i)}_{\rho_i'\rho_i}));$$

(ii) there is a regular Bohr set B' with $\mathrm{rk}(B') \leqslant k_i + \widetilde{O}(\alpha_i^{-4/3})$ and

$$\mu_G(B') \geqslant \exp(-\widetilde{O}(k_i + \alpha_i^{-4/3}))\mu_G(B^{(i)}_{\rho_i'\rho_i}) \geqslant \exp(-\widetilde{O}(k_i + \alpha_i^{-4/3}))\mu_G(B^{(i)})$$

such that $\|1_A * \beta'\|_{L^\infty(\mu_G)} = \widetilde{\Omega}(\alpha^{2/3})$;

(iii) there is a regular Bohr set $B^{(i+1)}$ with $k_{i+1} \leqslant k_i + \widetilde{O}(\alpha_i^{-2/3})$ and,

$$\mu_G(B^{(i+1)}) \geqslant \exp(-\widetilde{O}(k_i + \alpha_i^{-2/3}))\mu_G(B^{(i)}_{\rho_i'\rho_i}) \geqslant \exp(-\widetilde{O}(k_i + \alpha_i^{-2/3}))\mu_G(B^{(i)}),$$

such that $\alpha_{i+1} \geqslant \alpha_i(1 + c_{9.2}\alpha_i^{1/3}/4)$;

(iv) there is a regular Bohr set $B^{(i+1)}$ with $k_{i+1} \leqslant k_i + \widetilde{O}(\alpha_i^{-1})$ and,

$$\mu_G(B^{(i+1)}) \geqslant \exp(-\widetilde{O}(k_i + \alpha_i^{-1}))\mu_G(B^{(i)}_{\rho_i'\rho_i}) \geqslant \exp(-\widetilde{O}(k_i + \alpha_i^{-1}))\mu_G(B^{(i)}),$$

such that $\alpha_{i+1} \geqslant \alpha_i(1 + c_{9.2}/4)$.

In the second case we apply Theorem 8.1 to this translate and we are done:

$$T(1_A, 1_A, 1_A) \geqslant \exp(-\widetilde{O}(k_i + \alpha_i^{-4/3}))\mu_G(B^{(i)})^2.$$

Otherwise we can be in the third case at most $\widetilde{O}(\alpha^{-1/3})$ times and in the fourth case $O(\log 2\alpha^{-1})$ times. It follows that the iteration must terminate at some step $i_0 = \widetilde{O}(\alpha^{-1/3})$, where we shall have $k_{i_0} = \widetilde{O}(\alpha^{-1})$, and

$$\mu_G(B^{(i_0)}) \geqslant \exp(-\widetilde{O}(\alpha^{-4/3}))\mu_G(B),$$

from which the result follows. □

ACKNOWLEDGEMENT

The author should like to thank an anonymous referee for useful comments.

REFERENCES

[Beh46] F. A. Behrend. On sets of integers which contain no three terms in arithmetical progression. *Proc. Nat. Acad. Sci. U. S. A.*, 32:331–332, 1946.

[Bou85] J. Bourgain. Sidon sets and Riesz products. *Ann. Inst. Fourier (Grenoble)*, 35(1):137–148, 1985.

[Bou90] J. Bourgain. On arithmetic progressions in sums of sets of integers. In *A tribute to Paul Erdős*, pages 105–109. Cambridge Univ. Press, Cambridge, 1990.

[Bou99] J. Bourgain. On triples in arithmetic progression. *Geom. Funct. Anal.*, 9(5):968–984, 1999.

[Bou08] J. Bourgain. Roth's theorem on progressions revisited. *J. Anal. Math.*, 104:155–206, 2008.

[Cha02] M.-C. Chang. A polynomial bound in Freĭman's theorem. *Duke Math. J.*, 113(3):399–419, 2002.

[Elk08] M. Elkin. An improved construction of progression-free sets. arXiv:0801.4310, 2008.

[GK09] B. J. Green and S. V. Konyagin. On the Littlewood problem modulo a prime. *Canad. J. Math.*, 61(1):141–164, 2009.

[GR07] B. J. Green and I. Z. Ruzsa. Freĭman's theorem in an arbitrary abelian group. *J. Lond. Math. Soc. (2)*, 75(1):163–175, 2007.

[Gre03] B. J. Green. Some constructions in the inverse spectral theory of cyclic groups. *Combin. Probab. Comput.*, 12(2):127–138, 2003.

[GT08] B. J. Green and T. C. Tao. An inverse theorem for the Gowers $U^3(G)$ norm. *Proc. Edinb. Math. Soc. (2)*, 51(1):73–153, 2008.

[GW10a] W. T. Gowers and J. Wolf. Linear forms and quadratic uniformity for functions on $\mathbb{Z}_N$. arXiv:1002.2210, 2010.

[GW10b] B. J. Green and J. Wolf. A note on Elkin's improvement of Behrend's construction. In *Additive number theory: Festschrift in honor of the sixtieth birthday of Melvyn B. Nathanson*, pages 141–144. Springer-Verlag, 1st edition, 2010.

[HB87] D. R. Heath-Brown. Integer sets containing no arithmetic progressions. *J. London Math. Soc. (2)*, 35(3):385–394, 1987.

[Lev04] V. F. Lev. Progression-free sets in finite abelian groups. *J. Number Theory*, 104(1):162–169, 2004.

[Mes95] R. Meshulam. On subsets of finite abelian groups with no 3-term arithmetic progressions. *J. Combin. Theory Ser. A*, 71(1):168–172, 1995.

[Rot52] K. F. Roth. Sur quelques ensembles d'entiers. *C. R. Acad. Sci. Paris*, 234:388–390, 1952.

[Rot53] K. F. Roth. On certain sets of integers. *J. London Math. Soc.*, 28:104–109, 1953.

[Shk06] I. D. Shkredov. On sets of large trigonometric sums. *Dokl. Akad. Nauk*, 411(4):455–459, 2006.

[Shk07] I. D. Shkredov. Some examples of sets of large trigonometric sums. *Mat. Sb.*, 198(12):105–140, 2007.

[Shk08a] I. D. Shkredov. On sets of large trigonometric sums. *Izv. Ross. Akad. Nauk Ser. Mat.*, 72(1):161–182, 2008.

[Shk08b] I. D. Shkredov. On sets with small doubling. *Mat. Zametki*, 84(6):927–947, 2008.

[Shk09] I. D. Shkredov. On sumsets of dissociated sets. *Online J. Anal. Comb.*, (4):Art. 4, 26, 2009.

[SS42] R. Salem and D. C. Spencer. On sets of integers which contain no three terms in arithmetical progression. *Proc. Nat. Acad. Sci. U. S. A.*, 28:561–563, 1942.

[SY10] I.D. Shkredov and S. Yekhanin. Sets with large additive energy and symmetric sets. *Journal of Combinatorial Theory, Series A*, In Press, Corrected Proof:–, 2010.

[Sze90] E. Szemerédi. Integer sets containing no arithmetic progressions. *Acta Math. Hungar.*, 56(1-2):155–158, 1990.

[TV06] T. C. Tao and H. V. Vu. *Additive combinatorics*, volume 105 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2006.

Department of Pure Mathematics and Mathematical Statistics, University of Cambridge, Wilberforce Road, Cambridge CB3 0WB, England

E-mail address: `t.sanders@dpmms.cam.ac.uk`